

The OptlPuter: Global OptlPortals on Lambdas

Dr. Thomas A. DeFanti

Research Scientist

**California Institute for Telecommunications and Information Technology
University of California, San Diego**

**Distinguished Professor Emeritus of Computer Science
University of Illinois at Chicago**

Director

**Electronic Visualization Laboratory
University of Illinois at Chicago**

Principal Investigator, TransLight/StarLight



First, A Vision for the Next Decade: Gigapixels @ Terabits/sec

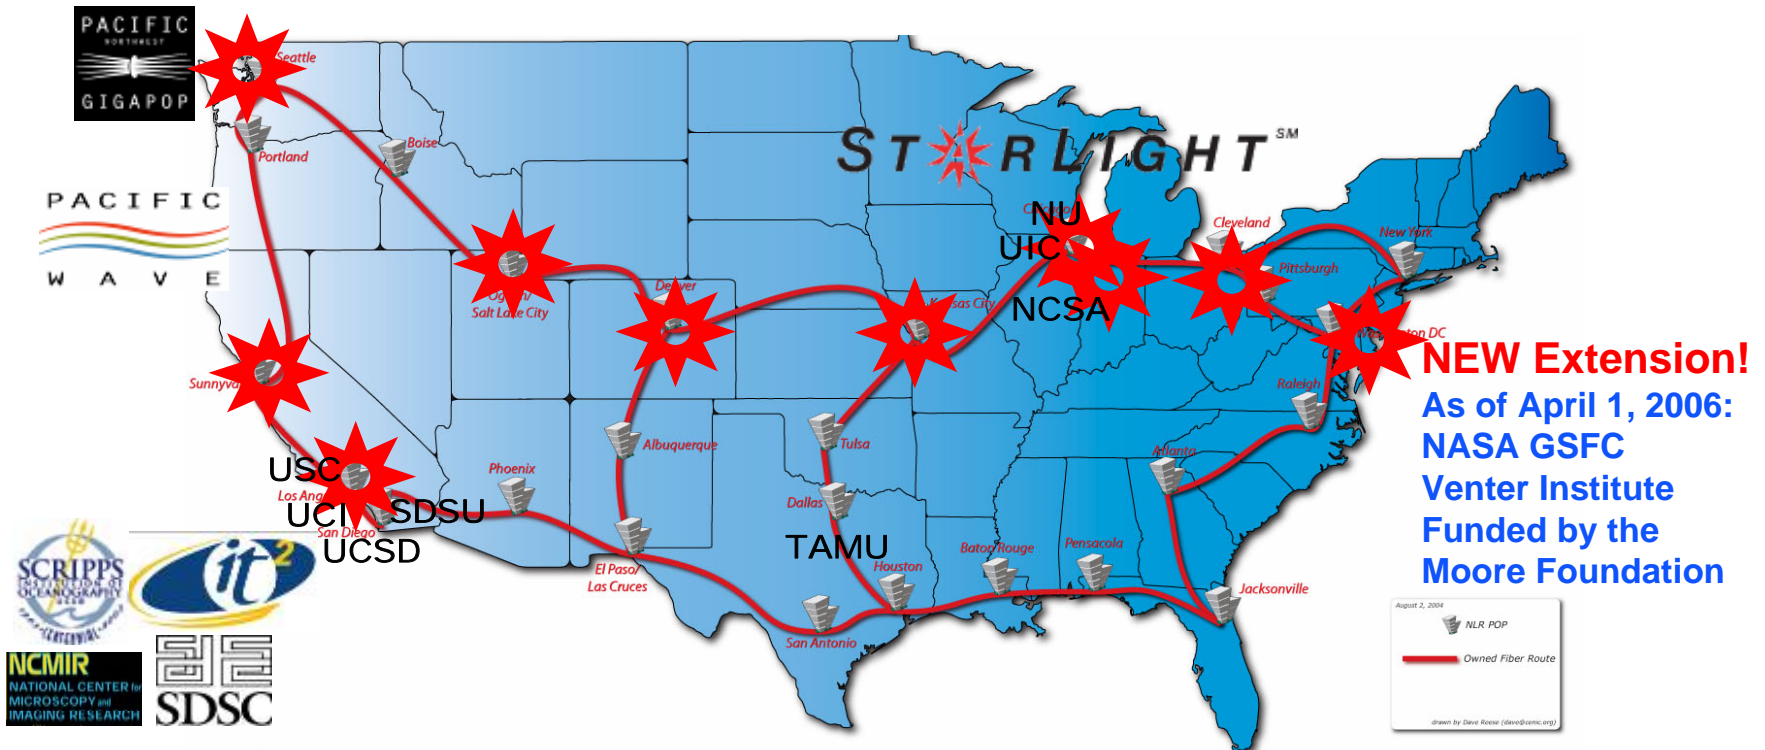
4K
Streaming Video

Gigapixel
Wall Paper

Augmented Reality
No Glasses

1 GigaPixel x 3 Bytes/pixel x 8 bits/byte x 30 frames/sec ~ 1 Terabit/sec!

Since 2004: The OptIPuter's 10GE CAVEwave on the National LambdaRail



 **CAVEwave™ is the University of Illinois at Chicago (UIC) Electronic Visualization Laboratory's very own 10 Gigabit wavelength on the NLR infrastructure, connected to the University of Washington in Seattle and UCSD in San Diego, enabling OptIPuter experiments.**

The CAVEwave links the California Institute for Telecommunications and Information Technology (Calit2)

- Virtual Reality, 4K Digital Cinema, HDTV
- Nanotech, Chips, Radio, Photonics, Grid Software, Data, Applications
- Over 1000 Researchers in Two Buildings
- Linked via Dedicated Optical Networks
- For International Conferences and Testbeds



UC Irvine




**Preparing for an World in Which
Distance Has Been Eliminated...**




UC San Diego

And the Electronic Visualization Laboratory at UIC

- EVL established in 1973
- EVL is the development laboratory for eruptive technology research and deployment for Calit2 and many other laboratories



9Gb/s Disk-to-Disk Performance with LambdaStream between EVL and Calit2



CAVEWave:

20 senders to 20 receivers (point to point)

Effective Throughput = 9.01 Gbps
(San Diego to Chicago)

450.5 Mbps disk to disk transfer per stream

Effective Throughput = 9.30 Gbps
(Chicago to San Diego)

465 Mbps disk to disk transfer per stream

TeraGrid:

20 senders to 20 receivers (point to point)

Effective Throughput = 9.02 Gbps
(San Diego to Chicago)

451 Mbps disk to disk transfer per stream

Effective Throughput = 9.22 Gbps
(Chicago to San Diego)

461 Mbps disk to disk transfer per stream


Dataset: 220GB Satellite Imagery of Chicago courtesy USGS.

Each file is 5000 x 5000 RGB image with a size of 75MB i.e ~ 3000 files

How We Get to Gigapixels over Terabits: The OptIPuter Project

- An NSF-funded ITR award to develop technology for the real-time collaboration and visualization of very-large, time-varying volumetric datasets for the earth sciences and the biosciences; now extending to metagenomics and digital cinema applications through private sector funding
- OptIPuter is our model of computing in which ultra-high-speed networks form the backplane of a global computer thereby removing bandwidth as an obstacle in data-intensive sciences
- UCSD, UIC, UCI, USC, SDSU, NU, UIUC/NCSA, TAMU, UvA, SARA, CANARIE, CRC, AIST, KISTI



NIH Biomedical Informatics Research Network



<http://ncmir.ucsd.edu/gallery.html>



NSF EarthScope and ORION



siovizcenter.ucsd.edu/library/gallery/shoot1/index.shtml



www.optiputer.net



OptIPuter System Software Architecture

Applications

Applications/ Web Services

Visualization

Telescience

Data Services

Visualization

SAGE

JuxtaView

LambdaRAM

Vol-a-Tile

Distributed Virtual Computer Middleware

DVC Runtime Library

DVC Services

DVC Job
Scheduling

DVC
Communication

DVC Core Services

Resource
Identify/Acquire

Namespace
Management

Security
Management

High Speed
Communication

Storage
Services

PIN/PDC

Globus

GRAM

GSI

XIO

RobuStore

Optical Signaling, Management

TRANS LIGHT



GTP

XCP

UDP

CEP





LambdaStream

High-Speed Transport Protocols

The OptlPortal

Enabling Users to See Gigabyte Data Objects

- 100s of megapixel 2D images & 3D objects are becoming common
- Usable interactive analysis and visualization of data objects requires deterministic, not “best effort” campus networks
 - Guaranteed Bandwidth (data movement, security)
 - Guaranteed Latency (visualization/collaboration, data analysis)
 - Guaranteed Availability (remote instruments, production schedules)
- Interactive analysis and visualization of high-res data objects requires:
 - Scalable visualization displays (OptlPortals)
 - Scalable adaptive graphics environments (SAGE)






OptIPortals are 40-200 Megapixel Collaboration Walls with 5-60 Terabytes Storage Connected by Lightpaths



Price: \$1K/Megapixel (10 pixels/JPY) (plus storage and network as needed)

Some Important OptIPuter OptIPortal Sites



SAGE OptIPortal Software: 10 Wireless Laptop Users All Pushing Their Desktops to the EVL OptIPortal--Goal is a Distributed Gigapixel in 2007



SAGE Also Merges Data with HDTV and 4K Streams



Emergency Response Prototype OptlPortal: Waterproof, Portable, Wireless



Calit2 @ UC Irvine Has the Largest Known OptIPortal—the HPerWall

**Calit2@UCI Apple Tiled Display Wall
Driven by 25 Dual-Processor G5s
50 Apple 30" Cinema Displays
200 Million Pixels of Viewing Real Estate!**

HDTV

**8 MP Digital Cameras
4K Digital Cinema**

**Data—One Foot Resolution
USGS Images of La Jolla, CA**



Source: Falko Kuester,
Calit2 UCI



Next: The OptlPortal in 3D, No Glasses: EVL's Varrier High-Resolution VR Autostereo Display




Collaboration is the OptIPuter's Main Contribution

- Many of the highest performance e-science applications involve national and international collaboration.
- This was NSF's main goal of building Chicago's STAR TAP (ATM) in 1996 and StarLight (GE and 10GE) in 2001.
- Colleagues in Japan, in America, Canada, Netherlands, Korea, China, UK, Czech Republic and elsewhere, agreed in 2003 to form a global initiative to create a photonic network testbed for the common good.
- This, of course, is *GLIF*, the Global Lambda Integrated Facility.
- The OptIPuter depends on GLIF and its GOLEs



The OptIPuter Jumps Beyond Traditional Web-Accessible Databases



This block contains three screenshots of web interfaces. The left screenshot shows the "A General Integrated View and Specimen" page from the BIRN database, featuring a hierarchical diagram of data types (Projects, Experiment, Subjects, Process) and a search form. The middle screenshot shows the "XPSSS (xml-based Protein Structure Search Service)" page from the PDB database, with a search form and a list of protein structures. The right screenshot shows the "Entrez, The Life Sciences Search Engine" page from NCBI Genbank, displaying a search bar and a list of search results.

BIRN

STARLIGHT

TRANS LIGHT



Source: Phil

Papadopoulos, SDSC,
Calit2

NCBI Genbank



OptIPuter's Direct Access Core Architecture Will Create Next Generation Metagenomics Server

Sargasso Sea Data

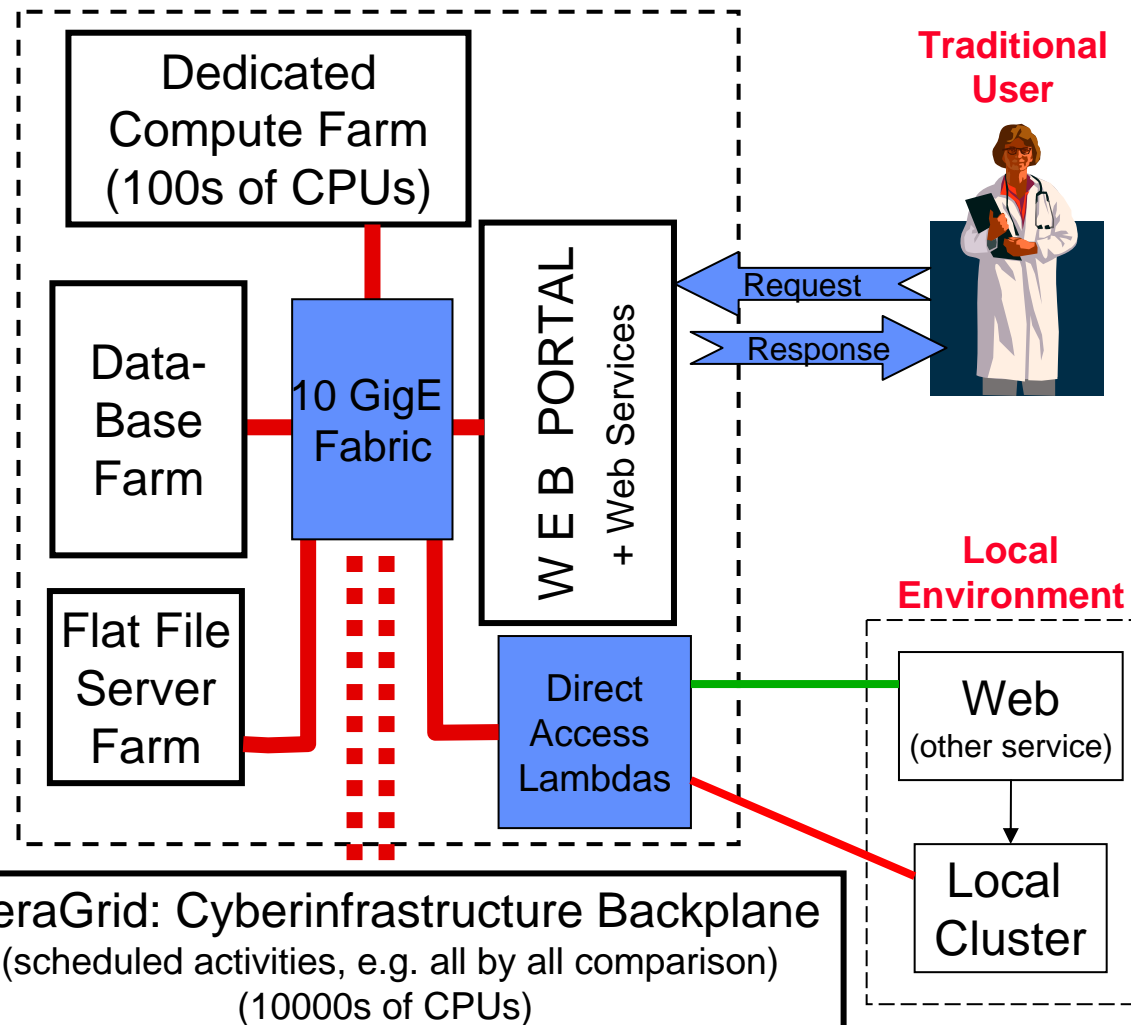
Sorcerer II Expedition
(GOS)

JGI Community
Sequencing Project

Moore Marine
Microbial Project

NASA Goddard
Satellite Data

Community Microbial
Metagenomics Data



J. Craig Venter
INSTITUTE

SDSC
SAN DIEGO SUPERCOMPUTER CENTER



Source: Phil Papadopoulos,
SDSC, Calit2




CAMERA: Community Cyberinfrastructure for Advanced Marine Microbial Ecology Research and Analysis

National LambdaRail
Direct Connect
Computation and Storage Complex

Funded by: Gordon and Betty Moore Foundation



Joint Partnership of:

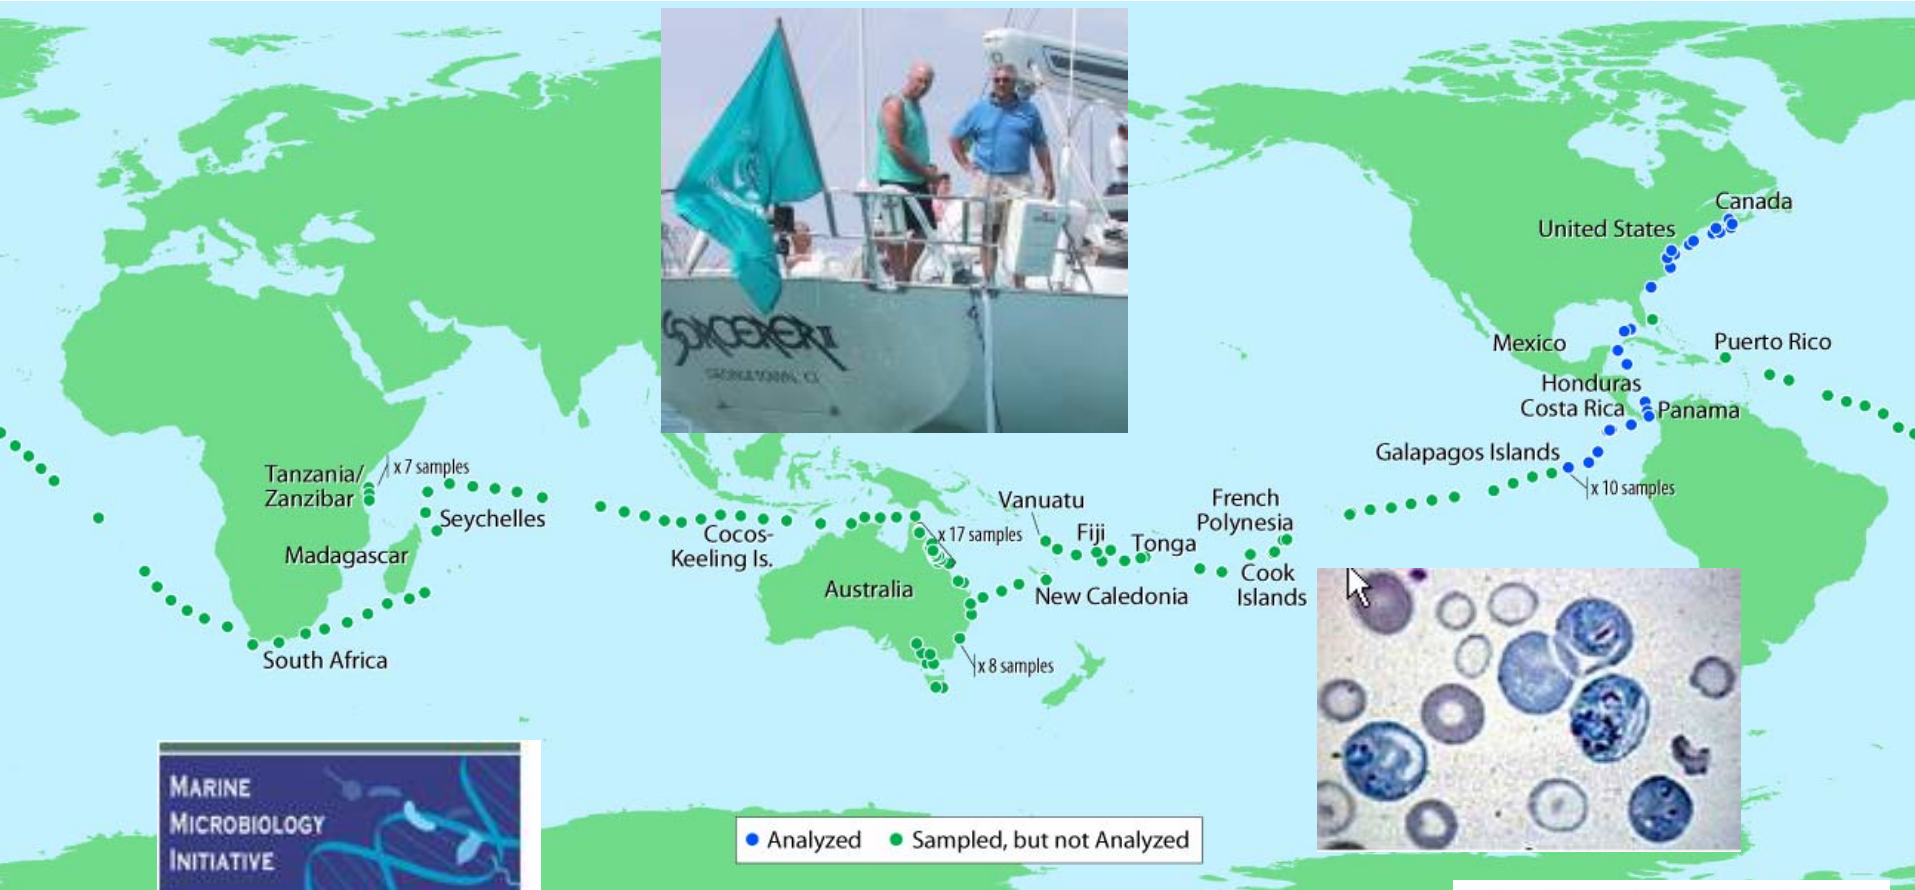


PI Larry Smarr



Marine Genome Sequencing Project

Measuring the Genetic Diversity of Ocean Microbes



**CAMERA will include
All Sorcerer II Metagenomic Data**

J. Craig Venter
INSTITUTE



Conclusion: OptIPuter Applications Illustrate Need for Optical Networks

- **Interactive collaboration using video (SD, HD, SHD) and/or VR**
 - Low latency streaming (real-time use)
 - High data rates
 - Lossy protocols OK
 - Multi-channel, multi-cast
- **Interactive Biomedical Imaging**
 - Very high resolution 2D (tens to hundreds of megapixels)
 - Volume visualizations (billions of zones in 3D)
- **Interactive Geoscience Imaging**
 - Very high resolution 2D (tens to hundreds of megapixels)
 - Volume visualizations (billions of zones in 3D)
- **Digital cinema**
 - Large data sets—4K is 250Mb/s compressed, up to 14Gb/s uncompressed
 - Security
- **Metagenomics**
 - Large computing
 - Large data sets

Conclusions after 4 Years of OptIPuter R&D

- OptIPuter technologies work with effort, but getting easier to configure, more affordable, and quite replicable
- OptIPuter applications benefit from deterministic networks:
 - Known and knowable bandwidth
 - Known and knowable latency
 - Availability of entire 10G lightpaths when necessary
- OptIPuter partner activities are training the next generation of network engineers, middleware, and application experts
- High-resolution video and audio are extremely demanding
- Scalability and commercialization remain challenging!

Conclusion: OptIPortal Gigapixel Terabit/sec Networked Data-Intensive Displays are Coming!

4K
Streaming Video

Gigapixel
Wall Paper

Augmented
Reality, no Glasses

Source: Jason Leigh, EVL

1 GigaPixel x 3 Bytes/pixel x 8 bits/byte x 30 frames/sec ~ 1 Terabit/sec!



Thank You Very Much!

- **OptIPuter/OptIPortal research, education, and outreach efforts are made possible, in major part, by funding from:**
 - US National Science Foundation (NSF) awards ANI-0225642, EIA-0115809, and SCI-0441094
 - State of Illinois I-WIRE Program, and major UIC cost sharing
 - State of California and UCSD Calit2
 - Many corporate friends and partners
 - Gordon and Betty Moore Foundation
- **Argonne National Laboratory and Northwestern University for StarLight and I-WIRE networking and management**



For More Information



- www.glif.is
- www.startap.net
- www.evl.uic.edu
- www.calit2.edu
- www.nlr.net